The Chatbot Will See You Now

Can chatbots be a supplementary element in doctor-patient communication?

Kjersti Torheim Bjelkarøy
Department of Design
Norwegian University of Science and Technology

ABSTRACT

General practitioners want patients to take more control of their own health. They also want to be guided and supported in the communication and diagnostic process with patients. This is something that can be accomplished with machine learning technology. When designing an interface of a health care system that is built on machine learning, patients' values in doctor-patient communication along with several human-computer interaction (HCI) design strategies becomes important guidelines. Chatbots are examples of an interface that builds on AI and machine learning technology. The design of a chatbot in a patient care context must allow for an understanding of how people desire to interact with it, in order to gain trust and acceptance. Although several strategies and principles can serve as guidelines on how to design a natural interaction between machine learning systems and humans, the interaction conventions should also be allowed to evolve naturally.

KEYWORDS: Machine learning, chatbot, doctor-patient communication, Human-computer interaction, artificial intelligence, user experience, health care.

1 INTRODUCTION

Machine learning technology can trace its roots back to the late 1950s, and the reader have probably encountered the technology at some point, either by using the Siri and Alexa voice assistants, Snapchats's facial recognition or the Netflix recommendations feature. Machine learning opens up for a whole new world of possibilities not only within the entertainment industry, but also in the industry of medicine. Machine learning is being used in image recognition algorithms that can discover tumour features relevant to prognosis. Mobile applications analyse the user's mental mood over time in order to provide coping strategies. Clinical support systems help doctors choose the right cancer treatment plan for an individual patient. (Marr, 2016)

From the point where a patient begins to experience symptoms until the point where he is diagnosed and starts a treatment, a complex service exists with several stakeholders and touchpoints. The patient's journey differs depending on the decease and symptoms, and the patient might be interacting with doctors, health personnel at the laboratory, hospitals, pharmacies, receiving tests results in the mail, seeking advice with his family or online, waiting for a referral to see a specialist, taking medication and following a treatment plan, the list goes on… Is it possible that machine

learning somewhere along the road might contribute to improve the patient journey?

Human communication and interaction is undoubtedly valued in primary health care service. A patient is assigned a doctor as "their" general practitioner and for many patients, this is the only "connection" they have to their own health. A rapidly growing trend within the field of machine learning are chatbots. Are chatbots entitled a role in doctor-patient connection? How can it affect the experience of going to the doctor's? What will it take for patients to trust in a virtual doctor as much as he trusts in a human doctor? What do patients value in doctor-patient communication and how can that be translated to a human-computer interaction? How can the interface between the technology and the user be designed so that the interests and values of both patient and doctor are preserved?

In order to answer the questions above, design strategies in human-computer interaction involving machine learning is explored in this paper. Further, what patients value in doctor-patient communication is discussed. This literature review seeks to purpose as a guideline for how human-computer interaction design principles should be married with what patients' value in doctor-patient communication, and how an interface that gathers patient information should be designed to build patient trust and acceptance.

2 BACKGROUND

2.1 What is a chatbot?

A chatbot system is defined as a "software program that simulate conversation with human users, using text, voice or images or a combination of spoken and visual heuristics"(Shawar & Atwell, 2007). Several terms are used to describe chatbots such as machine conversation systems, conversational agents, virtual agents and dialogue systems. Common for all is that they utilize a computer program that can perform tasks automatically or with minimal human intervention.

The use of chatbots have had a tremendous growth over the decade. It is predicted that by 2019, virtual personal assistants "will have changed the way users interact with devices and become universally accepted as part of everyday life" (Meulen, 2016). Klopfenstein, Delpriori, Malatini & Bogliolo (2017) draws parallels between the growth of chat bots and the growth of instant messaging (IM) applications. IM apps take up the top spots among the most used applications, and the attention is growing as their success depends on the network effect of their users. As conversing through messaging is becoming one of the most common ways humans communicate today, the conventions of messaging systems are becoming more and more familiar and recognizable .(Klopfenstein, Delpriori, Malatini, & Bogliolo, 2017). When chatbots are making use of these conventions, they are engaging with users in ways that they are already comfortable with.

Chatbot design employ different degrees of human-like appearance and behaviour, such as facial expressions, compassion, humour, tone of voice, mimicry and "chit-chatting". Etlinger (2017) suggests that in order to power an experience with chatbots, several guidelines should be taken into account, for example the ability to sense and respond appropriately to emotional signals from the user (Etlinger, 2017).

2.2 AI, machine learning and chatbots

It is important to separate between the terms artificial intelligence(AI), chatbots and machine learning. A chatbot is basically an interface

designed as a conversation, and is what meets the eye whereas AI and ML stays behind the stage curtain. Machine learning is a term used for the algorithms that gives systems the ability to learn from data, and is a subset of AI. Artificial intelligence, with a sci-fi ring to its name, is harder to define. In short, it is the branch of computer science that deals with the simulation of intelligent behaviour in computers. What exactly intelligent behaviour from a computer is, and what people perceive as intelligent behaviour from a computer, that is another question that this review paper will stir into.

2.3 Chatbots and primary care

The question arises: must a physician be human? Ellie, a computer developed at the Institute for Creative Technologies asks questions that a doctor might, such as "How easy is it for you to get a good night's sleep? Ellie analyses the patient's verbal responses, facial expressions and vocal intonations, and tries to detect signs of posttraumatic stress disorder or depression among others. In a randomized study, probands were told that Ellie was controlled by either a human or a computer program. Those who were told the latter ended up revealing more personal material to Ellie (Lucas, Gratch, King, & Morency, 2014).

Primary care services
Primary care services are commonly known as all health care services that do not involve institution or hospitals. Examples are GP (general practitioner) services and nursing services (Braut, 2009). In Norway, every person is assigned a GP through the municipality who works as patients' "first point of contact" in health related issues.

Pain problems in primary care services
From interviews conducted as research related to this review paper, it was discovered that several general practitioners have a wish for patients be more involved in, and to take better control of their own health. The GPs claimed that the information that patients give them the first 10 minutes of a consultation would be useful to possess prior to the consultation. This is when the patient tells the doctor about his symptoms and gives a short story of the background for his problems. The GPs suggested that by giving the patients a chance to "check-in" prior to a consultation, the information can be used to build a database that through machine learning technology can work as a decision support tool for doctors.

A common phenomenon among GPs are so-called "doorknob questions". Patients often come to a doctor's appointment with several issues that he wants to discuss, even though he might have informed about only one issue at the time of booking the appointment. Whether the patient decide to mention them, often depends if there is more time left of the appointment, and if the patients feel like bringing them up. Patients have the tendency to trivialize and neglect certain issues, even though they might be bothering them a lot. They postpone it until the very end of the consultation, and mention it to the doctor as they touch the doorknob. In many cases, these issues are too serious for the doctor to ignore.

Research methodology
This review paper aims to explore how knowledge about desired patient-doctor communication and strategies within human-computer interaction should affect the design of an interface that builds on machine learning technology. It is based on a literature review in addition to depth interviews of doctors. The goal of the interviews was to address pain problems in primary care services. The literature search covered terms like principles of human-computer interaction, doctor-patient

communication, chatbots, natural language processing, and user experience.

3 FINDINGS

3.1 Human-computer interaction styles

According to Shneiderman (1997) it exists five primary styles of which a human can interact with a computer: direct manipulation, menu selection, form fill-in, command language and natural language. Direct manipulation is a style where complex commands are replaced with visual or physical representations of the world of action, and the user is able to carry out tasks rapidly and observe the results immediately. Examples are dragging and dropping an element to the trash icon, or rotating the steering wheel in an automobile to the left, to turn left. Menu selection is when a user read a list of items, select one item and observe the effect. Form fill-in is when the user enters data in an input field and this interaction often involves a keyboard. Command language is when a user uses a certain syntax to give commands to a computer. Examples are keyboard shortcuts and programming. When these four styles were defined, the fifth one was described as a Star Trek scenario. Natural language is when a computer responds to commands that the user issue by speaking or typing in natural language. Shneiderman states that "the key impediment to natural language interaction is the habitability of the user interface - that is how easy it is for users to determine what objects and actions are appropriate", as natural language interaction interfaces provide less cues for the semantics of interaction than the four other interaction styles.

Natural language interaction (NLI) is the fundament for how a chatbot interact with the user. Natural language is also the fundament for how humans interact with each other, and when a system or a machine utilizes NLI and starts to learn, suggest and to take decisions, it is necessary for the acceptance of the system, that the user trust in it. Trust becomes a broad term in this context, and the following sections present research on how this trust can be cultivated and sustained. Without this trust, the process of gaining the data needed in machine learning algorithms, may stagnate.

3.2 Design strategies

Designing for transparency and disclosure
Users trust learning system more when they understand how the system arrives at its decisions. The user is also better able to correct and improve a system when the internals of its operation is visible and explained (Girardin, 2016). Stumpf and colleagues (2009) argues that a system's explanations of why it has made a prediction must be usable and useful to the user, and viable for processing by both the user and the system. (Stumpf et al., 2009)

A study reports that patient willingness to disclose information increases when a virtual human is used in clinical interviews instead of an actual human (Lucas et al., 2014). The study states that "compared to those who believed they were interacting with a human operator, participants who believed they were interacting with a computer reported lower fear of self-disclosure, lower impression management, displayed their sadness more intensely, and were rated by observers as more willing to disclose." Further, it states that "patients are particularly afraid to disclose personal, sensitive or stigmatizing information" and that patients try to selectively represent themselves in ways that they believe will make healthcare professionals view them positively.

Design for involvement, control and active learning

Computers have become indispensable parts of our daily lives and we have become increasingly accustomed to their reliability. As these systems become even more intelligent, the user's expectation of reliability will not lower (Borenstein, 2016). Reviewed as a good design practice, the user should be able to implicitly or explicitly tell about poor information. It is the designer's job to find ways to make this a part of the experience (Girardin, 2016).

Active learning is subfield of machine learning where the learning algorithm is allowed to choose the data from which it learns (Settles, 2010). The algorithm is "curios" and will eventually perform better with less training. A machine that "actively" learns is guided by an oracle (e.g a human annotator) that helps label training instances, and this way a person can teach a machine a task by showing examples.

Active learning has had success in accelerated learning in applications like text classification and object recognition(Amershi, 2015). Described as one of The Eight Golden Rules of interface design, supporting internal locus of control is important as users strongly desire the sense that they are in charge of the interface (Shneiderman, 1997). When accuracy matters, allowing the user to help could make a crucial difference (Stumpf et al., 2009). Stumpf et al. suggests that a rich human–computer collaboration is a promising direction for machine learning systems to work more intelligently, hand-in-hand with the user. Blythe (2005) argues that an obvious way to gather user feedback is to allow interactions in natural language (Blythe, 2005)

However, active learning is also proven to frustrate users. (Settles, 2010). Cakmak, Chao, & Thomaz (2010) discovered in a study where a human is teaching simple shapes to a robot and the robot requests particular examples, that "the constant stream of questions from the robot during the interaction was perceived as imbalanced and annoying", and that it led to a decline in the user's mental model of how the robot learned (Cakmak, Chao, & Thomaz, 2010). These studies reveals that systems need to account for human factors such as interruptibility or frustration. (Amershi, 2015)

Designing for conversation and tone of voice

Colby (1999) argues that "we need not take human-human conversation as the gold standard for conversational exchanges [with computers]. If one had a perfect simulation of a human conversant, then it would be human-human conversation and not human-computer conversation with its sometimes odd but pertinent properties" (Colby, 1999).

"Tone of voice" consists of both how a person speaks and the words he uses, and says a lot about how a person or a character is perceived. How factual, intelligent, concise, technical or informal a chatbot is experienced, can be designed though carefully choosing the chatbot's tone of voice. As chatting with machines is a relative new experience for many people, the design of the chat has to take into account that, yet, there "isn't really a social convention for accusing a person of being a bit of software." (Alasdair, 2017).

A website's tone of voice of any piece of content can be analysed along 4 dimensions: humour, formality, respectfulness, and enthusiasm (Meyer, 2016). In a study conducted by Nielsen Norman Group, two tone-of-voice samples called Epsilon and Zeta were evaluated in a hospital patient care context. Epsilon's tone profile is serious, formal, respectful and matter-of-fact. Zeta's tone profile is serious, casual, respectful and enthusiastic. Nielsen Norman Group found that

Zeta scored higher on the user perception of friendliness, trustworthiness and desirability (Meyer, 2016).

Designing for new modes of input

Machine learning enables deeper forms of communication with computers, and this requires that designers rethink some of the longstanding principles of user interface, user experience design and human-computer interaction. Human-computer interaction is no longer restricted to mouse and keyboard, but expand from visual inputs, aural inputs, corporeal inputs and environmental inputs (Hebron, 2016).

Zadrozny (2000) argues that the best way to facilitate human-computer interaction (HCI) is by allowing users "to express their interest, wishes, or queries directly and naturally, by speaking, typing, and pointing (Zadrozny et al., 2000). Negroponte (1970) adds that "it is gestures, smiles, and frowns that turn a conversation into a dialogue" (Negroponte, 1970).

Designing to meet expectations

In general, the aim of chatbot designers should be to build tools that help people facilitate their work and their interaction with computers using natural language; but not to replace the human role totally, or imitate human conversation perfectly (Shawar & Atwell, 2007). Etlinger (2017) argues that "bots that do not disclose at the outset that they are not human can frustrate and alienate users (Etlinger, 2017) and Tunkelang (2016) states that "chatbots should promote their strengths while exposing their weaknesses" (Tunkelang, 2016).

Chatbots with low human-like appearance are the ones most commonly used today. Examples span from Slack's planner bot, to Apple's Siri. Siri is designed as a coloured electronic wave, giving the user little indication of what it might have looked like, if it had a body. Chatbots that take form in more human-looking ways may be described as virtual humans(VS). VS are defined as "robots that are embodied, or animated characters that interact with people in a natural way" (Lucas et al., 2014).

Separating between these two types by the degree of human-like appearance may help us understand the concept of the uncanny valley. The concept was proposed by the Japanese roboticist Masahiro Mori in a 1970 essay, where he envisioned people's reactions to robots that looked and acted like humans. He hypothesized that "a person's response to a humanlike robot would abruptly shift from empathy to revulsion as it approached, but failed to attain, a lifelike appearance" (Mori, MacDorman, & Kageki, 2012). Tunkelang (2016) argues that the risk of falling into this valley is not just determined by the fact that chatbots are not lifelike enough, but by the fact that they do not provide the right affordances. Users "are not able to easily figure out what they can and cannot do".

Doctor-patient communication

According to Ong , de Haes, Hoos & Lammes (1995), doctor-patient communication has three different purposes: (a) creating a good inter-personal relationship between the physician and the patient, (b) exchanging information between the doctor and the patient, and (c) to make decisions related to medical treatment (Ong, de Haes, Hoos, & Lammes, 1995). Further, Ong et al. categorize patient outcomes as (1) patient satisfaction, (2) patient adherence to treatment, (3) understanding of information, and (4) health outcomes. Studies shows that good doctor-patient communication can lead to improved adherence (Haskard Zolnierek & DiMatteo, 2009), and improved health outcomes (Stewart, 1995).

Smith, Polis & Hadac (1981) argues that the medical interview length is positively correlated to patient satisfaction, and while an increased time by doctors reviewing the medical chart has a negative impact on patient satisfaction (Smith, Polis, & Hadac, 1981). Doctors need information from patients in order to diagnose them correctly and to create a treatment plan. Time pressure and high workload affects this information gathering, and often patients are interrupted when describing their problems to the doctor. This premature interruption of patients may lead to a potential loss of relevant information (Beckman & Frankel, 1984).

Ong et al. (1995) states that patients need to know why they are experiencing the problems they are, and what significance it has for them. This strong need for information is often underestimated by doctors. In addition, patients often have difficulties understanding and remembering what the physician has told them about a diagnosis and the prognosis of the disease (Ley, 1988). In a study from primary care, patients were only able to recall between 40 and 80% of the information they received (Ley, 1988).

Although the pathway from good doctor-patient dialogue to positive health care outcomes is not always clear, an example of a direct pathway can be seen when the dialogue with the doctor has a therapeutic effect on the patient (Alsos, 2011). Eye contact is important in doctor-patient communication as it enhances listening skills and makes physicians more effective in reading emotional cues (Roter & Hall, 2006).

4 DISCUSSION

In this section the findings from the research is analysed and discussed in the light of how it should relate to the design of a chatbot.

4.1 Implementing the design strategies for a chatbot in health care

Transparency and disclosure

When a doctor suggests a diagnose to a patient, the patient will expect a reasoning behind the doctor's suggestion or else he may decide to not believe in it and hence not follow the treatment. Since a machine learning system can take a larger amount of medical factors into account than a doctor can, it becomes even more important to make the line of reasoning of the machine visible to the patient.

How accurately an algorithm predicts a value can be measured, and this is where the machine may receive more trust than the doctor. By presenting to the user how likely the predicted value is to be correct, the system is admitting that it is not flawless and that it is up to the user to decide whether he wishes to trust in it and follow the recommendations or not.

When patients selectively choose what information they give the doctor with the agenda of having the doctor view them positively, information that can be important for them to disclose to healthcare professionals is in risk of being lost. With this in mind, an interesting question arises: will patients be less scared of being negatively judged by a chatbot, and hence give a better clinical interview? If a patient talks to a chatbot prior to a consultation and is asked to name all the problems he has in mind, could chatbots be a supportive tool to eliminate door knob questions?

Involvement, control and active learning

A patient's medical record is currently only written by doctors and other health care professionals. In Norway, patients are able to view their medical record, but not everyone knows about this and it is not a straight forward process. As long as patients are kept

out of the documentation process, they will not feel that they are obliged to take part. Hence, all the responsibility of writing a correct medical record falls on the doctor. Activating the patients by having a chatbot ask them to give feedback to the system, can improve patient engagement and trust, and prevent patient passiveness when it comes to diagnosing and documentation.

Tone of voice and conversation

So how should the tone of voice of a chatbot in a health care context be? The Golden Rule "do to others what you want them to do to you" also applies to machine learning. In customer services the customer is often polite to the service people, seeking the outcome of being judged as a nice person that deserves good service. People reflect their personality through the way they talk and act, and this way other people get an impression of who we are. Today, as instant messaging becomes a big part of our communication, we also reflect our personality through smileys and emoji. Will the chatbot be programmed so that it is able to judge us based on how we express us? If not, why should we be polite towards it? Why would we use a conversational tone with someone, or something, that will not judge us if we rather go straight to the point?

Designing for new modes of input

Allowing new modes of input like camera and microphone opens up for new possibilities, from interpreting facial expressions and vocal sounds of patients with depression, to diagnosing skin deceases from photographs. The question is, should we design for a future where we expect that computers can hear and see as good as humans can? If not, in what cases should the machine be excused, and when can we expect it to be even better than humans?

Designing to meet expectations

As chatbots are engaging and communicating to users using the similar messaging interfaces people use to communicate to each other, a clear expectation clarification needs to be set between the user and the chatbot. What the chatbot can help the patient with and what receiving that help will require from the patient will have to be clearly specified from day one. There is no doubt that a machine in a diagnosing context can have a larger and more complex image "in the back of its mind" than what a doctor may have, but only if the patient agrees to share and specify the data needed for building that image. For the patient to be willing to do this, there has to be a clear communication of what is in it for the patient.

It is also important to keep it crystal clear at all times if the patient is talking to another human, or to a machine. If the conversation opens up for involving both doctor and chatbot, the design must clearly communicate to whom the patient is speaking to, and it must be clear what the doctor and what the chatbot can do for the patient respectively, and what the information they request will be used for.

Doctor-patient communication

As stated above, interrupting patients during a consultation may result in a potential loss of relevant information. This is a good example of how a chatbot has the ability to work as a supplementary factor in doctor-patient communication. The patient may talk to the chatbot prior to a consultation, helping the chatbot build an overview of the situation the patient is in. This information can then be presented to the doctor in an understandable and straightforward way, reducing the time spent on reviewing patient charts during the consultation. In addition, a chatbot is not affected by time pressure and high workload when it comes to gathering patient information, and will not interrupt the patients

when they are describing their problems. This, along with the theory that patients are more willing to disclose when they are interacting with a computer may present support that a medical interview is better when supplemented with a chatbot.

As patients may struggle to remember the information they are given during a consultation, a chatbot could work as a helping assistant for the patient, summarizing and preserving the plan the doctor either passed orally or written, and remind the patient of it. It could for example be a reminder of when he should take his medicine, and also prompt the patient with an evaluation of the treatment he was given. If parts of the screening interview takes place in a chat, the results will be kept as a written conversation for the patient to read later if needed.

When health care professionals underestimate the importance of passing information to the patient, the chatbot may assist. The chatbot can give the patient the option to choose how much explanation they would like. The chatbot also has the advantage that it will always remember to explain it in everyday language, and not only in a medical language, as doctors often do.

The power of the human-human interaction of exchanging eye contact is hard to translate into human-computer interaction. Even though the chatbot can aid the doctor in getting more information from the patient, it must not be forgotten that the patient also often has the need and health benefits of talking to another human.

4.2 A side note from a user experience designer

How people experience interacting with machine learning technology has everything to do with what they expect from interacting with it. These expectations and goals are what user experience designers specialize in. User experience designers needs to be present in the development and utilization of machine learning technology in order to remind the industry that to successfully harness the possibilities that machine learning opens up for, it needs to be done in a way that is desired by the end user.

5 CONCLUSION

The design strategies in this paper are viewed as important when the interaction between a human and a computer is happening in a patient care context. Other principles that can be further researched are patient vulnerability and data security, design for user forgiveness, design for new modes of output (hardware) and design for uncertainty.

This paper has presented and discussed the term natural language interaction, and showed several references to how human-computer interaction can be empowered if humans can communicate using their natural ways of interacting. It has also presented arguments for how human-human conversation should not be seen as a gold standard when designing a human-computer conversation. The conclusion this paper wants to draw, is that when machine learning technology is involved, natural language interaction between human and computer is not the same as “natural” interaction between a human and a computer. The industry is trying to pre-define conventions for how people desire to naturally interact with machine learning technology interfaces, when really this is something that must evolve over time. The best practice business strategy is to be prepared for what the “survival of the fittest” conventions are.

A chatbot can successfully be designed to work as a helping tool in doctor-patient

communication, but it must be emphasized that it should work as a supplement and never a replacement. The chatbot and machine learning industry is too young to start replacing human interaction, especially if we are designing for good experiences. We are still at a place where we in all cases are able to tell the difference between a human and a robot, but we are also stepping into an era where we are interacting with computers the same way we interact with each other. The most important strategy when designing for patient trust is meeting their expectations and being clear about the affordances of the chatbot. In this way, a natural interaction will evolve over time.

When creating solutions that depend on data that learns from human behaviours, the process of gathering that behaviour data needs to follow design principles that facilitates motivation, engagement and trust from the user. It is up to user experience designers, in partnership with developers, psychologist, adventuresome anthropologists and sociologists to discover how this should be done.

6 REFERENCES

Alasdair, A. (2017). Augmented Reality Will Take Off As Soon As it Doesn't Look Silly. Retrieved from https://motherboard.vice.com/en_us/article/gv7ygb/augmented-reality-will-take-off-as-soon-as-it-doesnt-look-silly

Alsos, O. A. (2011). *Mobile Point of-Care Systems in Hospitals: Designing for the Doctor-Patient Dialogue.* Doktoravhandlinger ved NTNU (273).

Amershi, S., Cakmak, Maya, Knox, W. Bradley, Kulesza, Todd. (2015). Power to the People: The Role of Humans in Interactive Machine Learning. Retrieved from http://bradknox.net/public/papers/AIMagazine2015-IML.pdf

Beckman, H. B., & Frankel, R. M. (1984). The effect of physician behavior on the collection of data. *Ann Intern Med, 101*(5), 692-696.

Blythe, J. (2005). *Task learning by instruction in tailor*. Paper presented at the Proceedings of the 10th international conference on Intelligent user interfaces, San Diego, California, USA.

Borenstein, G. (2016). Power to the People: How One Unknown Group of Researchers Holds the Key to Using AI to Solve Real Human Problems. Retrieved from https://medium.com/@atduskgreg/power-to-the-people-how-one-unknown-group-of-researchers-holds-the-key-to-using-ai-to-solve-real-cc9e75b1f334

Braut, G. S. (2009). Primærhelsetjenesten. Retrieved from https://snl.no/prim%C3%A6rhelsetjeneste

Cakmak, M., Chao, C., & Thomaz, A. L. (2010). Designing Interactions for Robot Active Learners. *IEEE Transactions on Autonomous Mental Development, 2*(2), 108-118. doi:10.1109/TAMD.2010.2051030

Colby, K. M. (1999). Comments on Human-Computer Conversation. In Y. Wilks (Ed.), *Machine Conversations* (pp. 5-8). Boston, MA: Springer US.

Etlinger, S. (2017). *The Conversational business: How chatbots will reshape digital experiences*. Retrieved from One Bush Street CA 94104: file:///Users/kjerstibjelkaroy/Downloads/ConversationalBusiness_FINAL.pdf

Girardin, F. (2016). Experience Design in the Machine Learning Era. Retrieved from https://medium.com/@girardin/experience-design-in-the-machine-learning-era-e16c87f4f2e2

Haskard Zolnierek, K. B., & DiMatteo, M. R. (2009). Physician Communication and Patient Adherence to Treatment: A Meta-analysis.

Medical care, 47(8), 826-834. doi:10.1097/MLR.0b013e31819a5acc

Hebron, P. (2016). *Machine Learning for Designers*. 1005 Gravenstein Highway North, Sebastopol: O'Reilly Media.

Klopfenstein, L. C., Delpriori, S., Malatini, S., & Bogliolo, A. (2017). *The Rise of Bots: A Survey of Conversational Interfaces, Patterns, and Paradigms*. Paper presented at the Proceedings of the 2017 Conference on Designing Interactive Systems, Edinburgh, United Kingdom.

Ley, P. (1988). *Communicating with patients: Improving communication, satisfaction and compliance*. New York, NY, US: Croom Helm.

Lucas, G. M., Gratch, J., King, A., & Morency, L.-P. (2014). It's only a computer: Virtual humans increase willingness to disclose. *Computers in Human Behavior, 37*(Supplement C), 94-100. doi:https://doi.org/10.1016/j.chb.2014.04.043

Marr, B. (2016). How Machine Learning, Big Data And AI Are Changing Healthcare Forever. Retrieved from https://www.forbes.com/sites/bernardmarr/2016/09/23/how-machine-learning-big-data-and-ai-are-changing-healthcare-forever/ -70f924e11a1c

Meulen, R. (2016). Gartner Says by 2019, 20 Percent of User Interactions With Smartphones Will Take Place via VPAs [Press release]. Retrieved from https://www.gartner.com/newsroom/id/3551217

Meyer, K. (2016). The Four Dimensions of Tone of Voice. Retrieved from https://www.nngroup.com/articles/tone-voice-samples/

Mori, M., MacDorman, K. F., & Kageki, N. (2012). The Uncanny Valley [From the Field]. *IEEE Robotics & Automation Magazine, 19*(2), 98-100. doi:10.1109/MRA.2012.2192811

Negroponte, N. (1970). *The Architecture Machine: Toward a More Human Environment*: MIT Press.

Ong, L. M., de Haes, J. C., Hoos, A. M., & Lammes, F. B. (1995). Doctor-patient communication: a review of the literature. *Soc Sci Med, 40*(7), 903-918.

Roter, D., & Hall, J. A. (2006). *Doctors talking with patients/patients talking with doctors : improving communication in medical visits* (2nd ed ed.): Praeger.

Settles, B. (2010). *Active Learning Literature Survey*. Retrieved from http://burrsettles.com/pub/settles.activelearning.pdf

Shawar, B. A., & Atwell, E. (2007). Chatbots: Are they Really Useful? *LDV Forum, 22*, 29-49.

Shneiderman, B. (1997). *Designing the User Interface: Strategies for Effective Human-Computer Interaction*: Addison-Wesley Longman Publishing Co., Inc.

Smith, C. K., Polis, E., & Hadac, R. R. (1981). Characteristics of the initial medical interview associated with patient satisfaction and understanding. *J Fam Pract, 12*(2), 283-288.

Stewart, M. A. (1995). Effective physician-patient communication and health outcomes: a review. *Cmaj, 152*(9), 1423-1433.

Stumpf, S., Rajaram, V., Li, L., Wong, W.-K., Burnett, M., Dietterich, T., . . . Herlocker, J. (2009). Interacting meaningfully with machine learning systems: Three experiments. *International Journal of Human-Computer Studies, 67*(8), 639-662. doi:https://doi.org/10.1016/j.ijhcs.2009.03.004

Tunkelang, D. (2016). The Uncanny Valley of Chatbots. Retrieved from https://medium.com/@dtunkelang/the-uncanny-valley-of-chatbots-9f265b42d83b

Zadrozny, W., Budzikowska, M., Chai, J., Kambhatla, N., Levesque, S., & Nicolov, N. (2000). Natural language dialogue for personalized interaction. *Commun. ACM, 43*(8), 116-120. doi:10.1145/345124.345164